arXiv:1411.5715v4 [math.ST] 6 Aug 2015

Weak continuity of predictive distribution for Markov survival processes

Walter Dempsey and Peter McCullagh *

August 10, 2015

Abstract

We explore the concept of a consistent exchangeable survival process - a joint distribution of survival times in which the risk set evolves as a continuous-time Markov process with homogeneous transition rates. We show a correspondence with the de Finetti approach of constructing an exchangeable survival process by generating iid survival times conditional on a completely independent hazard measure. We describe several specific processes, showing how the number of blocks of tied failure times grows asymptotically with the number of individuals in each case. In particular, we show that the set of Markov survival processes with weakly continuous predictive distributions can be characterized by a two-dimensional family called the harmonic process. We end by applying these methods to data, showing how they can be easily extended to handle censoring.

1 Introduction

1.1 Background

This paper introduces the concept of a consistent exchangeable survival process, which is a joint distribution p_n of survival times $T_1, \ldots, T_n$ such that the temporal progress of the risk set is Markovian with homogeneous transition rates. The approach is similar to that taken by Kingman (1980) or Aldous (1996) in constructing coalescent-like trees, or by McCullagh, Pitman and Winkel (2006) in characterizing self-similar Gibbs fragmentation trees. Each consistent survival process is generated by a splitting rule which determines the joint distribution for fixed n and the predictive distributions, $\mathrm{pr}(T_{n+1} \in A \,|\, T_1, \ldots, T_n)$. We show that each splitting rule is defined by a corresponding characteristic index, $\zeta = \{\zeta_i\}_{i=1}^{\infty}$. In particular, the marginal distributions are exponential with rate ζ_1.

Three closely related processes in the class are explored in detail, the power process for which $\zeta_n \propto n^{\rho}$ for $0 < \rho < 1$, the gamma process for which $\zeta_n \propto$

*Department of Statistics, University of Chicago, 5734 University Ave, Chicago, Il 60637, U.S.A.

$\log(1+n/\rho)$ and the harmonic process for which $\zeta_n \propto \sum_{j=0}^{n-1} 1/(\rho+j)$. The latter has a simple form for the joint density, it is easy to generate sequentially, and the conditional distributions have a close affinity with various long-standing estimators for the survival distribution, such as the Kaplan-Meier estimator (Kaplan and Meier, 1958). Moreover, the harmonic process is shown to be the *only* Markov survival process with weakly continuous predictive distributions.

The de Finetti approach to constructing an exchangeable survival process is to generate survival times conditionally independent and identically distributed given a completely independent hazard measure, i.e. the cumulative conditional hazard is a Lévy process (Cornfield and Detre, 1977; Kalbfleisch, 1978; Hjort, 1990; Clayton, 1991). Such processes are sometimes called *neutral to the right* (Doksum, 1974; James, 2006). Each process in this class is determined by the characteristic exponent of the Lévy measure, and the temporal evolution of the risk set is Markovian, i.e., they are exchangeable Markov survival processes. Thus, the two approaches, which are mathematically natural in very different ways, give rise to the same set of exchangeable processes. The probability density can be computed using the characteristic index, and the process can be simulated directly from the predictive distributions, by-passing the Lévy process entirely.

A survival process in this class typically has multiple tied failure times, and these ties generate an exchangeable partial ranking, or an ordered partition, of the particles. The partial ranking is infinitely exchangeable, and can be generated by a sequential assignment rule similar to the Chinese restaurant process for the Dirichlet partition (Ferguson, 1973, 1974; Pitman, 2006). As a function of n, the number of blocks may be bounded, or it may grow at a logarithmic or super-logarithmic rate, or it may grow at rate n^α for some $0 < \alpha \leq 1$. For the harmonic and the gamma processes, the rate is asymptotically proportional to $\log^2 n$.

1.2 Risk set trajectory

Consider a set of patients or particles in which particle i survives to time $T_i > 0$, not necessarily finite. To each particle there corresponds an observation interval $[0, c_i]$, and the event time recorded for particle i is $Y_i = \min(T_i, c_i)$. In all of the mathematical theory that follows, the censoring times are arbitrary positive numbers and are known for each particle, so if $Y_i = c_i$ the event is known to be a censoring time; otherwise, if $Y_i < c_i$, the event is known to be a death or failure. Evidently, setting $c_i = 0$ is equivalent to removing particle i from the sample.

It is helpful mathematically to think of the observation on particle i not as a real number but as a survival interval

$$\begin{aligned} R_i &= [0, Y_i) \quad \text{if} \quad Y_i < c_i \\ R_i &= [0, Y_i] \quad \text{if} \quad Y_i = c_i. \end{aligned}$$

Alternatively, and equivalently, the observation is a Boolean function such that

$R_i(t) = 1$ if $t \in R_i$ and zero otherwise. With this notation

$$R(t) = \{i : R_i(t) = 1\}$$

is the subset of particles known to be alive at time t, i.e., the risk set at time t. The temporal trajectory of R is a non-increasing sequence of subsets beginning at $R(0) = [n] = \{1, \ldots, n\}$, which is the sample of particles initially under observation. The right increment $R(t) \setminus R(t^+)$ is the subset of particles censored at time t; the left increment $B(t) = R(t^-) \setminus R(t)$ is the subset of particles observed to fail at time t. It is worth emphasizing that the subsets $B(t)$ and $R(t)$ are disjoint for every t.

The number of individuals in the risk set $R^\sharp(t) = \#R(t)$ is a non-increasing real-valued function, beginning at $R^\sharp(0) = n$, decreasing in integral steps at each censoring and failure time to $R^\sharp(t) = 0$ for $t > \max(Y)$. The right increment $R^\sharp(t) - R^\sharp(t^+)$ counts censored particles; the left increment $R^\sharp(t^-) - R^\sharp(t)$ counts failures, and is regarded as a right-continuous counting measure.

This paper is concerned with survival processes, by which is meant probabilistic models for the temporal evolution of the risk set, i.e., stochastic models in continuous time with state space equal to the subsets of $[n]$. Since n is arbitrary, each stochastic process must be defined in a consistent manner for general samples in such a way that p_n is the marginal distribution of p_{n+1} under deletion, either by setting $c_{n+1} = 0$ or by removal from the sample. Censoring is not ignored, but the discussion presumes $c_i = \infty$ to avoid unnecessary distractions. Although it affects the evolution of the risk set, censoring has a relatively trivial effect on all probability calculations and associated statistical procedures such as parameter estimation and the calculation of conditional distributions such as the survival function $\mathrm{pr}(T_{n+1} > t \,|\, R[n])$ or the residual survival function $\mathrm{pr}(T_{n+1} > t \,|\, R[n], T_{n+1} > s)$, where $R[n]$ denotes the risk set trajectory restricted to the set of particles, $[n]$. The natural convention on left and right increments is consistent with standard practice (Kalbfleisch, 1978, Table 1), and is sufficient for procedures to be expressed in the generality required to accommodate arbitrary fixed right censoring.

2 Exchangeable Markov survival processes

2.1 Partial rankings

A partial ranking of $[n]$ is an *ordered list* $B = (B_1, \ldots, B_k)$ consisting of disjoint non-empty subsets of $[n]$ whose union is $[n]$. The terms *partial ranking* and *ordered partition* are used interchangeably. The elements of $[n]$ are unordered within blocks, but B_1 is the subset ranked first, B_2 is the subset ranked second, and so on. The alternative notation $B = B_1|B_2|\cdots|B_k$ is sometimes used, either for ordered or unordered partitions.

There are three partial rankings of [2], 12 of [3], and 75 of [4], arranged in equivalence classes (with respect to permutation of elements and permutation

of unequal-sized blocks) as follows;

$$\begin{array}{l}\mathcal{OP}_2 : 12; \quad 1|2,\ 2|1\\ \mathcal{OP}_3 : 123; \quad 12|3, 13|2, 23|1,\ 1|23, 2|13, 3|12; \quad 1|2|3(3!)\\ \mathcal{OP}_4 : 1234; \quad 123|4(4\times 2); \quad 12|34(6); \quad 12|3|4(12\times 3); \quad 1|2|3|4(4!)\end{array}$$

Thus $12|3|4(12\times 3)$ means that there are three equivalence classes of types $211, 121, 112$ depending on whether the first, second or third block contains two elements, and that each class contains 12 partial rankings generated by permutation of elements.

To each unordered partition of $[n]$ containing k blocks, there correspond $k!$ partial rankings, one for each ordering of the blocks. Thus, if $S_{n,k}$ is number of partitions of $[n]$ containing k blocks, i.e., Stirling's number of the second kind, then

$$\#\mathcal{P}_n = \sum_{k=1}^{n} S_{n,k}, \qquad \#\mathcal{OP}_n = \sum_{k=1}^{n} k!\, S_{nk}$$

is the number of partitions of $[n]$ and the number of partial rankings of $[n]$ respectively. The numerical values for $n \le 10$ are as follows:

n	1	2	3	4	5	6	7	8	9	10
$\#\mathcal{P}_n$	1	2	5	15	52	203	877	4 140	21 147	115 975
$\#\mathcal{OP}_n$	1	3	13	75	541	4 683	47 293	545 835	7 087 261	102 247 563

2.2 Exchangeable Markov partial rankings

A random partial ranking of $[n]$ is a probability distribution $p_n(\cdot)$ on the finite set $\mathcal{OP}_n$. The distribution is *finitely exchangeable* if $p_n(\cdot)$ depends only on the block sizes and the block order: in general $p_n(B_1, \ldots, B_k) \neq p_n(B_{\sigma(1)}, \ldots, B_{\sigma(k)})$ for a permutation σ of the blocks.

Definition 2.1 (Markovian)**.** *The distribution p_n is said to be* Markovian *if for each subset of size $1 \le m \le n$ there exists a splitting distribution $q(s,d)$ with $s + d = m$ and $d \ge 1$ such that the partial ranking $B = (B_1, \ldots, B_k)$ of $[n]$ consisting of k blocks of sizes $B_1^\sharp, \ldots, B_k^\sharp$ has probability*

$$\begin{aligned} p_n(B) &= q(n - B_1^\sharp, B_1^\sharp) \times p_{n-B_1^\sharp}(B_2, \ldots, B_k) \\ &= \prod_{j=1}^{k} q(n - B_1^\sharp - \cdots - B_j^\sharp, B_j^\sharp). \end{aligned}$$

The following proposition is a direct consequence of the above definition.

Proposition 2.2. *Every non-negative function, $q(r,s) > 0$ for $s > 0$, subject to normalization*

$$\sum_{d=1}^{n} \binom{n}{d} q(n-d, d) = 1 \tag{2.1}$$

for all $n \ge 0$ defines a Markovian survival process.

In the context of survival models, the first argument of q is the number of survivors, and the second is the number of failures at each successive hazard, so $q(s,d) \neq q(d,s)$ is not symmetric in its arguments. In this sense, the splitting distribution for partial rankings is very different from the splitting distribution for Markov fragmentation trees (Aldous 1996; McCullagh, Pitman and Winkel 2006). The normalization conditions are also different.

2.3 Kolmogorov consistency

A sequence $p = (p_1, p_2, \ldots)$ in which p_n is a probability distribution on partial rankings of $[n]$ is said to be Kolmogorov-consistent if each p_n is the marginal distribution of p_{n+1} under the map $\mathcal{OP}_{n+1} \to \mathcal{OP}_n$ in which the element $n+1$ is ignored or deleted. This condition implies that the censoring time for one particle, c_{n+1}, for example, has no effect on the survival times for other particles. The concern here is the conditions for consistency of exchangeable Markov partial rankings, so the choice of element to be deleted is immaterial. Without consistency, the sequence $\{p_n\}$ does not determine a process, there is no associated partial ranking of the infinite set (population), and there is no possibility of inference using conditional distributions.

Consider the family of Markov partial rankings generated by a splitting distribution $q(\cdot,\cdot)$. Proposition 2.3 details when such a splitting distribution corresponds with a Kolmogorov-consistent Markov survival process.

Proposition 2.3 (Consistency condition)**.** *The splitting distribution corresponds to a Kolmogorov-consistent Markov survival process if and only if*

$$(1 - q(n,1))q(n-d,d) = q(n-d,d+1) + q(n-d+1,d) \tag{2.2}$$

for all integers $n \geq d \geq 1$. In particular, equation (2.2) implies that the splitting probabilities are determined by the sequence of singleton splits $q(n,1)$ for $n \geq 0$.

Proof. To derive the conditions for consistency proceed by induction supposing that $p_1, \ldots, p_n$ are mutually consistent in the Kolmogorov sense. Then, in order for p_{n+1} to be consistent with p_n, it must be for each ordered partition B of $[n]$,

$$\begin{aligned} p_n(B) = q(n,1)p_n(B) &+ q(n - B_1^\sharp, B_1^\sharp + 1)p_{n-B_1^\sharp}(B_2, \ldots) \\ &+ q(n - B_1^\sharp + 1, B_1^\sharp)p_{n-B_1^\sharp+1}(\cdots). \end{aligned}$$

The terms appearing on the right are all of the events in $\mathcal{OP}_{n+1}$ such that deletion of $n+1$ gives rise to the ordered partition $B \in \mathcal{OP}_n$. Either $n+1$ occurs in the first block as a singleton, which has probability $q(n,1)p_n(B)$; or it occurs appended to B_1 as a non-singleton, which has probability $q(n - B_1^\sharp + 1, B_1^\sharp + 1)p_{n-B_1^\sharp}(B_2, \ldots)$; or it occurs elsewhere either as a singleton or appended to one of the blocks of B, which occurs with probability $q(n - B_1^\sharp + 1, B^\sharp)p_{n-B_1^\sharp+1}(\cdots)$. But, by the induction hypothesis that $p_1, \ldots, p_n$ are mutually consistent, the latter event occurs with probability $p_{n-B_1^\sharp}(B_2, \ldots)$. Hence, an exchangeable

Markov family of distributions is Kolmogorov-consistent if and only if, for each $n \geq 1$ and each ordered partition B of $[n]$,

$$(1 - q(n,1))q(n - B_1^\sharp, B_1^\sharp) = q(n - B_1^\sharp, B_1^\sharp + 1) + q(n - B_1^\sharp + 1, B^\sharp).$$

which is equivalent to equation (2.2). □

If the splitting rule corresponds to a Kolmogorov-consistent Markov survival process, it is said to be a *consistent splitting rule*. A nice property of such rules is that they admit simple integral representations.

Proposition 2.4. *Every consistent splitting rule admits an integral representation*

$$q(n-d,d) = q_n(d) = \frac{1}{Z_n}\left(\int_0^1 x^{n-d}(1-x)^d \varpi(dx) + c \cdot \mathbf{1}_{\{d=1\}}\right) \tag{2.3}$$

where $Z_n = \int (1-x^n)\varpi(dx) + n \cdot c$. *Here the measure is defined on* $[0,1)$ *and satisfies*

$$\int_0^1 (1-x)\,\varpi(dx) < \infty.$$

The constant, c, is called the erosion coefficient, and ϖ*, is called the dislocation measure.*

Both the condition and integral representation are similar to (4) in McCullagh, Pitman and Winkel (2008) and to proposition 41 of Ford (2006). However, the splitting probabilities in these papers are symmetric and subject to a normalization condition different from (2.1), so the probabilities are different.

2.4 The characteristic index

Here the *characteristic index* associated with each consistent exchangeable partial ranking is defined and a relation to the splitting rule via kth order differences is shown.

Definition 2.5. A *characteristic index*, ζ, is a sequence $\zeta_0, \zeta_1, \zeta_2, \ldots$ beginning with $\zeta_0 = 0$, $\zeta_1 > 0$, and subsequently

$$\zeta_{n+1} = \frac{\zeta_n}{1 - q_{n,1}} = \zeta_1 \prod_{j=1}^{n} (1 - q_{j,1})^{-1} \tag{2.4}$$

for $n \geq 1$.

The sequence is said to be in *standardized form* if $\zeta_1 = 1$, and each standardized sequence determines a Markov partial ranking provided that the splitting probabilities in (2.2) are non-negative.

A natural question is whether the splitting rules can be reconstructed from the characteristic index. Proposition 2.6 provides an affirmative answer via kth order forward differences, $\Delta^k\zeta$, defined as

$$(\Delta^k\zeta)_n = \sum_{j=0}^{k}(-1)^{k-j}\binom{k}{j}\zeta_{n+j},$$

for $n \geq 0$, so that $\Delta\zeta_n = \zeta_{n+1} - \zeta_n$ is the first difference, $(\Delta^2\zeta)_n = \zeta_{n+2} - 2\zeta_{n+1} + \zeta_n$ is the second, and so on.

Proposition 2.6. *Given a characteristic index, ζ, the corresponding splitting rule is given by*

$$q(r,d) = \frac{(-1)^{d-1}(\Delta^d\zeta)_r}{\zeta_{r+d}} \tag{2.5}$$

for $r \geq 0$ and $d \geq 1$. The splitting probabilities are therefore determined by the standardized sequence ζ/ζ_1. Moreover, the characteristic index yields a consistent splitting rule if it satisfies a non-negativity condition on forward differences; that is, $(-1)^{d-1}(\Delta^d\zeta)_r \geq 0$.

The proof of Proposition 2.6 can be found in Appendix B. The main advantage of working with the characteristic index is that the Kolmogorov consistency condition (2.2) may be written in a more convenient form as a non-negativity condition on forward differences.

2.4.1 Examples

Since each Markov partial ranking is associated with a characteristic sequence (modulo scalar multiplication), the space of Markov partial rankings may be identified with the space of non-negative sequences satisfying (2.5). Evidently this space is a convex cone, closed under positive linear combinations. It contains the trivial sequence $\zeta_n = n$, for which $\Delta\zeta = (1, 1, \ldots)$ and $\Delta^d\zeta \equiv 0$ for $d \geq 2$. This means $q(r,1) = 1/(r+1)$, so all splits are singletons, and there are no ties in the partial ranking.

For $\rho \geq 0$, the singleton split $q(n-1,1) = 1/(n+\rho)$ for $n \geq 2$ implies $\zeta_n \propto n+\rho$, linear in n for $n \geq 1$, with $\zeta_0 = 0$. The forward difference sequences imply that $\Delta\zeta_{n-1}/\zeta_n = 1/(n+\rho)$ for $n \geq 2$, $q(0,n) = (-1)^{n-1}(\Delta^n\zeta)_0 = \rho/(n+\rho)$ and $q(s,d) = 0$ otherwise. At each event time, one individual out of n fails with probability $1/(n+\rho)$ (for each individual); otherwise, with probability $\rho/(n+\rho)$, the entire population fails.

It also contains the sequences $\zeta_n = n^\alpha$ for $0 < \alpha < 1$, for which $q(r,1) = 1 - (1+1/r)^{-\alpha}$. Similarly, $q(n-1,1) = 1/n^2$ gives rise to $\zeta_n \propto n/(n+1)$, and the uniform splitting rule $q(n-d,d) = (n-d)!\,d!/n\,n!$.

For each $0 < \alpha < 1$, the non-decreasing geometric sequence $\xi_n = 1 - \alpha^n$ is such that

$$(-1)^{d-1}(\Delta^d\xi)_r = (-1)^d\alpha^r(\alpha-1)^d = \alpha^r(1-\alpha)^d > 0,$$

implying that there exists a Markov partial ranking with binomial splitting probabilities $q(r,d) = \alpha^r(1-\alpha)^d/\xi_{r+d}$ for $d \geq 1$. Moreover, for any non-negative measure w on $[0,1)$, the integrated sequence $\zeta_n = \int_0^1 (1-\alpha^n)\, dw(\alpha)$ is monotone, and is such that

$$(-1)^{d-1}(\Delta^d \zeta)_r = \int_0^1 \alpha^r (1-\alpha)^d \, dw(\alpha) \geq 0.$$

Thus, the Kolmogorov condition is automatically satisfied for all measures w such that $\zeta_1 = \int_0^1 (1-\alpha)\, dw(\alpha)$ is positive and finite. By equation (2.3), this is the complete set of splitting rules with zero erosion measure ($c = 0$). Finally, for each $\beta > 0$, the sequence $\zeta_{\beta n} = \int_0^1 (1-\alpha^{\beta n})\, dw(\alpha)$ also satisfies the Kolmogorov condition, so each ζ having an integral representation may be extended to a function on the positive real line.

Example 1 (Beta-splitting rules)**.** *A natural family of measures are those conjugate to the binomial splitting probabilities, $dw(\alpha) = \alpha^{\rho-1}(1-\alpha)^{\beta-1}d\alpha$. Here, $\rho > 0$ while $\beta > -1$ as $d \in \{1, \ldots, n\}$. The characteristic index then depends on the second parameter, β.*

1. *For $\beta > 0$, the characteristic index is*

$$\zeta_n = B(\rho,\beta) - B(\rho+n,\beta) = B(\rho,\beta)\left(1 - \frac{\rho^{\uparrow n}}{(\rho+\beta)^{\uparrow n}}\right) \tag{2.6}$$

   where $\rho^{\uparrow k} = \rho\cdot(\rho+1)\cdots(\rho+k-1)$ is the ascending factorial and $B(\rho,\beta)$ is the beta function. The previous mentioned example, $q(n-1,1) = 1/n^2$, corresponds to $\rho = \beta = 1$.

2. *For $\beta = 0$, the characteristic index is*

$$\begin{aligned}\zeta_n &= \int_0^1 \alpha^{\rho-1}(1+\alpha+\cdots+\alpha^{n-1})\, d\alpha \\ &= \sum_{j=0}^{n-1} \frac{1}{\rho+j} = \psi(n+\rho) - \psi(\rho),\end{aligned} \tag{2.7}$$

   where ψ is the derivative of the log gamma function. This process is called the harmonic process*.*

3. *Finally, for $\beta \in (-1,0)$, the process has characteristic index*

$$\begin{aligned}\zeta_n &= \sum_{j=0}^{n-1} \frac{\Gamma(j+\rho)\Gamma(\beta+1)}{\Gamma(j+1+\beta+\rho)} \\ &= O\left(n^{-\beta}\right)\end{aligned}$$

A very similar process to the harmonic process with density measure $dw(\alpha) = \alpha^{\rho-1}d\alpha/(-\log(\alpha))$ has characteristic index

$$\begin{aligned}\zeta_n &= \int_0^1 (1-\alpha^n)\alpha^{\rho-1}\,d\alpha/(-\log\alpha) = \int_0^\infty (1-e^{-nz})z^{-1}e^{-\rho z}\,dz \\ &= \log(1+n/\rho). \end{aligned} \tag{2.8}$$

This is the characteristic index of the gamma process, which is explored in more detail in section 3.

2.5 Holding times

A partial ranking of $[n]$ is not only an ordered partition $(B_1, \ldots, B_k)$, but also a strictly decreasing sequence of subsets

$$[n] \equiv R_0 \supset R_1 \supset \cdots \supset R_k \equiv \emptyset,$$

called risk sets, in which the increments $B_r = R_{r-1} \setminus R_r$ are non-empty. By associating with each risk set an independent exponentially distributed holding time, a Markov process in continuous time is constructed whose trajectories are non-increasing subsets of $[n]$, with the empty set as the absorbing state. Moreover, by choosing the rate function in an appropriate way, the survival process can be made both consistent under subsampling, and exchangeable under permutation of particles. Consistency under subset selection means that there exists an infinite process, i.e., a survival process for the infinite population, which implies in turn that the ratio p_{n+1}/p_n determines the conditional distribution of T_{n+1} given the survival times $T[n]$ for the initial sample. It is also possible to compute the conditional distribution given $T[n]$ and $T_{n+1} > 5$, i.e., to prognosticate in a mathematically consistent manner.

An argument essentially the same as that used in section 4 of McCullagh, Pitman and Winkel (2008) leads to the following consistency condition on the rate function

Proposition 2.7. *A sequence of rate functions, $\{\tau_n\}_{n=1}^\infty$, is consistent if*

$$\tau_{n+1}(1-q(n,1)) = \tau_n.$$

In other words, the characteristic index $\tau_n \equiv \zeta_n$ is also the exponential failure rate needed to ensure consistency of the continuous-time Markov process.

While the assumption of temporal homogeneity seems natural for the Markov model, the argument can be extended to non-constant holding rates. That is, we may specify a marginal hazard function, $h_2 : \mathbb{R}^+ \to \mathbb{R}$. Consistency implies the above relation holds with $h_n(t)$ replacing τ_n. Alternatively, a monotone transformation of time can be applied in order to modify the marginal distribution.

2.6 Density function

Since the evolution of the process $R(t)$ is Markovian, it is a straightforward exercise to give an expression for the probability density function at any specific temporal trajectory. The observation space consists of a partial ranking B of $[n]$ comprising $k = \#B$ disjoint subsets, and for each subset a failure time. The probability that the first failure occurs in the interval dt_1 and that B_1 is the set of failures is

$$\zeta(n)e^{-\zeta(n)t_1}\,dt_1 \times q(R_1^\sharp, B_1^\sharp) = e^{-\zeta(n)t_1}\,dt_1 \times \lambda(R_1^\sharp, B_1^\sharp),$$

where $\zeta(n) \equiv \zeta_n$ and $\lambda(r,d) = q(r,d)\zeta(r+d) = (-1)^{d-1}(\Delta^d\zeta)(r)$. Continuing in this way, it can be seen that the joint density at any temporal trajectory $R(\cdot)$ consisting of k blocks with failure times $0 < t_1 < \cdots < t_k$ is

$$f_n(B,t) = \exp\Big(-\int_0^\infty \zeta(R^\sharp(s))\,ds\Big)\prod_{j=1}^k \lambda\big(R^\sharp(t_j), B^\sharp(t_j)\big). \tag{2.9}$$

Here $k = \#B$ is the number of blocks, or more generally the number of distinct failure times, and $B_j \equiv B(t_j)$ is the block or subset of particles failing at time t_j. The density is non-negative, and the integral is one. In the absence of censoring, this means

$$\sum_{B\in\mathcal{OP}_n}\int_{0<t_1<\cdots<t_{\#B}} f_n(B,t)\,dt = 1,$$

so the number of blocks $1 \le \#B \le n$ is a random variable whose distribution is determined by (2.9), and hence by ζ. Supplementary figure 1 shows simulated survival processes with characteristic index $\zeta_n = n^\rho$ for various values of ρ together with the associated conditional survival function. Also, supplementary figure 2 shows several simulated harmonic survival processes for various choices of $\rho > 0$ along with the conditional survival function.

The random sequence of failure times $T_1, T_2, \ldots$ whose finite-dimensional joint distributions are given by (2.9) is infinitely exchangeable. The n-dimensional joint distribution is continuous in the sense that it has no fixed atoms. For $n \ge 2$, it is not continuous with respect to Lebesgue measure in $\mathbb{R}^n$ because the distribution has condensations on all diagonals implying that $\mathrm{pr}(T_1 = T_2) = q(0,2) > 0$, and likewise for arbitrary subsets. The one-dimensional marginal distributions are exponential with rate $\zeta(1)$. However, a monotone continuous temporal transformation that sends Lebesgue measure to the measure $\nu(\cdot)$ on $(0,\infty)$, also transforms (2.9) to an exchangeable semi-Markov process with density

$$f_n(B,t) = \exp\Big(-\int_0^\infty \zeta(R^\sharp(s))\,d\nu(s)\Big)\prod_{j=1}^k \lambda\big(R^\sharp(t_j), B^\sharp(t_j)\big)\,\nu(dt_j). \tag{2.10}$$

If $g(T) = \nu((0,T))$ is the associated monotone continuous function then $g^{-1}(T_i)$ is exponential with rate $\zeta(1)$.

Although the argument leading to (2.9) did not explicitly consider censoring, the density function has been expressed in integral form so that censoring is accommodated correctly. The pattern of censoring affects the evolution of $R^\sharp$, and thus affects the integral, but the product involves only failures and failure times.

2.7 Sequential description

Since f_n is the marginal distribution of f_{n+1}, the conditional distribution of T_{n+1} given the temporal evolution of the risk set $R[n] \equiv (B, t)$ for the first n particles is given by the ratio $f_{n+1}(B', t')/f_n(B, t)$, where (B', t') is any event compatible with the observation (B, t). If the new particle fails at one of the previous failure times, then $t' = t$, and B' is obtained by inserting the new particle into one of the blocks of B; otherwise the new particle fails interstitially as a singleton in one of the intervals, $(0, t_1), \ldots, (t_{\#B}, \infty)$, in which case $\#B' = \#B + 1$.

The conditional distribution is best described in terms of the conditional hazard measure, Λ, which is such that

$$\mathrm{pr}(T_{n+1} > t \,|\, R[n]) = e^{-\Lambda((0,t])}.$$

The conditional hazard has a continuous component supplemented by an atom at each previously observed failure time. The continuous component has a density and a cumulative hazard

$$\begin{aligned} h(t) &= \zeta(R^\sharp(t) + 1) - \zeta(R^\sharp(t)) = (\Delta\zeta)(R^\sharp(t)), \\ H(t) &= \int_0^t (\Delta\zeta)(R^\sharp(s))\, ds. \end{aligned}$$

Note that $R^\sharp$ is piecewise constant, so the integral is trivial to compute, but censoring implies that it is not necessarily constant between successive failures. For all consistent Markov survival processes, $\zeta_1 > 0$ implies that the continuous component has infinite total mass, so $\mathrm{pr}(T_{n+1} < \infty \,|\, R[n]) = 1$, i.e., T_{n+1} is finite with probability one.

Let t be a failure time with $R^\sharp(t^-) = r + d$ and $R^\sharp(t) = r$ with $d > 0$. At each such point the conditional hazard has an atom with finite mass

$$\Lambda(\{t\}) = \log \frac{\zeta(r + d)\, q(r, d)}{\zeta(r + d + 1)\, q(r + 1, d)},$$

or, on the probability scale,

$$\exp(-\Lambda(\{t\})) = \frac{\zeta(r + d + 1)\, q(r + 1, d)}{\zeta(r + d)\, q(r, d)} = \frac{(\Delta^d\zeta)(r + 1)}{(\Delta^d\zeta)(r)}.$$

The total mass of the atoms is finite.

Given the trajectory of the risk set for the first n particles, the conditional survival function is

$$\mathrm{pr}(T_{n+1} > t \,|\, R[n]) = \exp(-H(t)) \prod_{j : t_j \le t} \frac{(\Delta^{d_j}\zeta)(r_j + 1)}{(\Delta^{d_j}\zeta)(r_j)}. \tag{2.11}$$

Although this may look a little complicated, it is not difficult to generate the survival times sequentially for processes whose characteristic index admits a simple expression for finite differences. Right censoring is automatically accommodated by the integral in the continuous component, so the observed trajectory $R[n]$ may be incomplete.

The harmonic process (2.7) with characteristic index $\zeta_n = \nu(\psi(n+\rho)-\psi(\rho))$ for some $\nu > 0$ is such that $(-1)^{d-1}(\Delta^d\zeta)_r = \nu\Gamma(d)/(r+\rho)^{\uparrow d}$. Accordingly, the continuous component of the conditional hazard is $h(t) = \nu/(R^\sharp(t)+\rho)$, implying that

$$H(t) = \sum_{i:t_i \le t} \nu \frac{t_i - t_{i-1}}{R^\sharp(t_{i-1}) + \rho} + \nu \frac{t - t_j}{R^\sharp(t_j) + \rho},$$

where the sum runs over event times, censored or failure, such that $t_i \le t$, and t_j is the last such event. The discrete component (2.11) is a product over failure times

$$\prod_{j:t_j \le t} \frac{(\Delta^{d_j}\zeta)(r_j+1)}{(\Delta^{d_j}\zeta)(r_j)} = \prod_{j:t_j \le t} \frac{r_j + \rho}{r_j + d_j + \rho}. \tag{2.12}$$

In most cases of practical interest, the continuous component is negligible over the greater part of the range of interest; for small ρ, the discrete component is essentially the same as the right-continuous version of the Kaplan-Meier product limit estimator. Note that exchangeability implies the joint conditional survival probability, $\mathrm{pr}(T_{n+1} > t, T_{n+2} > t' \,|\, R[n])$, is distinct from the Kaplan-Meier product estimate which assumes independence among individuals.

2.8 Weak continuity of predictions

The conditional distribution of T_{n+1} given the sequence $\mathbf{t} = (t_1, \ldots, t_n)$ of previous failures has an atom at each distinct failure time, and is continuous elsewhere. The predictive distribution is weakly continuous if a small perturbation of the failure times gives rise to a small perturbation of the predictive distribution. In other words, for each $n \ge 1$, and for each non-negative vector $\mathbf{t}$,

$$\lim_{\epsilon \to 0} \mathrm{pr}(T_{n+1} \le x \,|\, T[n] = \mathbf{t} + \epsilon) = \mathrm{pr}(T_{n+1} \le x \,|\, T[n] = \mathbf{t})$$

at each continuity point, i.e., $x > 0$ not equal to one of the failure times.

Weak continuity is an additivity condition on the hazard atom as a function of the risk set and the number of tied failures. If d failures are exactly tied, the hazard atom is

$$-\log \frac{(\Delta^d\zeta)(r+1)}{(\Delta^d\zeta)(r)};$$

the same failures occurring as singletons one millisecond apart give rise to d distinct hazard atoms whose sum is

$$-\log \frac{(\Delta\zeta)(r+d)}{(\Delta\zeta)(r)}.$$

The predictive distribution is weakly continuous if these are equal for every integer $r \ge 0$ and $d \ge 1$.

Theorem 2.8 (Continuity of predictions)**.** *A Markov survival process has weakly continuous predictive distributions if and only if it is a harmonic process:* $\zeta(n) \propto \psi(n+\rho)-\psi(\rho)$ *for some* $\rho > 0$*. The iid exponential model is included as a limit point.*

To see that it is not satisfied by any other Markov survival process, it is sufficient to consider a sequence in standard form beginning with $(\Delta\zeta)_0 = 1$, $(\Delta\zeta)_1 = \rho/(\rho+1) < 1$. Then the key continuity condition determines the subsequent sequence $(\Delta\zeta)_r = \rho/(\rho+r)$ in conformity with the harmonic series. The only exception is the iid exponential process, which arises in the limit $\rho \to \infty$ in which tied failures occur with probability zero. All other Markov survival processes, including the gamma process, have predictive distributions that are discontinuous as a function of the initial configuration.

2.9 Self-similarity and lack-of-memory

Every Markov survival process has the property that the conditional joint distribution of the residual lifetimes $T_1 - t, \ldots, T_n - t$ given that $\min(T_1, \ldots, T_n) > t$ is the same as the unconditional distribution of $T_1, \ldots, T_n$. This property, called *lack of memory*, follows from the Markov property of the distribution (2.9).

In general, if $t > 0$ is fixed, only a subset of the initial sample of n particles will survive beyond that time. Let $S = \{i\colon T_i > t\} \subset [n]$ be the survivors. Every Markov survival process with consistent finite-dimensional distributions $\{p_n\}$ also has the property that the conditional joint distribution given S of the residual lifetimes $\{T_i - t\colon i \in S\}$ is distributed as $p_{S^\sharp}$. This property, called *self-similarity*, is a consequence of Markov homogeneity, namely that the transition intensities are consistent and constant in time.

Lack of memory is a property of the distribution p_n alone, and does not require consistency of p_n with p_{n+1}. By contrast, self-similarity is a property of the process. Obviously, self-similarity implies lack of memory.

2.10 Seeded series & urn models

Suppose that an initial sequence consisting of m values $T_1, \ldots, T_m$ is given, and that these correspond to an initial risk-set trajectory R_0. Subsequent values are generated using the Markov rule (2.11). The sequence $T_{m+1}, \ldots$ is called a seeded series because the joint distribution depends strongly on the initial sequence. It is natural to ask what the joint distribution of $T_{m+1}, \ldots, T_{m+n}$ is—whether it is stationary, whether it is exchangeable, and so on.

It is not difficult to see that the joint density of the $m+n$-configuration R given R_0 is

$$\exp\Big(-\int_0^\infty \big(\zeta(R^\sharp(s)) - \zeta(R_0^\sharp(s)\big)\, ds\Big) \prod_{\text{deaths}} \frac{\lambda(R^\sharp(t_j), B^\sharp(t_j))}{\lambda(R_0^\sharp(t_j), B_0^\sharp(t_j))}$$

where the zero-order difference is defined as $\lambda(r,0) = \Delta^0\zeta(r) = 1$. The product runs over all death times.

In other words, regardless of the initial configuration, the subsequent series is infinitely exchangeable. Although it is Markovian, the initial series introduces persistent temporal features, so the failure rate is not temporally homogeneous. The one-dimensional marginal distributions satisfy (2.11) (i.e. the distribution of T_i given R_0).

Supplementary figure 2 illustrates the impact of seeding on conditional survival function estimates. Specifically, the 50 initial values are assumed independent, uniform on $(1, 2)$. An additional 400 survival times are generated conditional on the seeded series given the harmonic process for a given choice of parameters ν and ρ.

2.11 Number of blocks & block sizes

The distribution of the number of blocks in a random ordered partition depends critically on the splitting probabilities. The mean number of blocks satisfies the recurrence relation

$$\mu_n = 1 + \sum_{d=1}^{n} \binom{n}{d} q(n-d, d)\mu_{n-d}$$

and there is a similar recurrence relation

$$M_n(t) = e^t \sum_{d=1}^{n} \binom{n}{d} q(n-d, d)M_{n-d}(t)$$

for the moment generating functions. The interest here is in the behaviour for large n.

At one extreme, $q(n-1, 1) = 1/n$ implies $q(r, d) = 0$ for $d > 1$, and $\mu_n = n$, so every block is a singleton and the partition is trivial. Similarly, $q(n-1, 1) = 1/(n+\rho)$ implies that the ratio μ_n/n has a strictly positive limit $1/(1+\rho)$, implying that a positive fraction of the blocks are singletons.

As for block sizes, it is clear that the distribution of the size of block i, $\#B_i$, stochastically dominates the distribution for all subsequent blocks. By a stick-splitting argument, the expected block size for any block can be derived by first examing $E[\#B_1]$. For $q(n-1, 1) = 1/(n+\rho)$, the asymptotic expected size of the first block is $(1+\rho)$.

The beta-splitting rules exhibit different behavior depending on the value of β. Lemma 2.9 provides a description of the asymptotic behavior of the number of blocks and block size for this process.

Lemma 2.9. *Let $D_n(\rho, \beta)$ denote the number of blocks given n individuals for the beta process with $\rho \in (0, \infty)$ and $\beta \in (-1, \infty)$. Then*

1. *For $\beta > 0$, as $n \to \infty$,*

$$E[D_n(\rho, \beta)] \sim \log(n)$$

where $x_n \sim y_n$ if $\lim_{n\to\infty} x_n/y_n = C \in (0, \infty)$ for x_n and y_n two non-negative sequences.

The fraction of edges in the first block, $n^{-1}\#B_{1,n}$, is asymptotically distributed $Beta(\beta, \rho)$. Therefore, the relative frequencies within each block are given by

$$(P_1, P_2, \ldots) = (W_1, \bar{W}_1 W_2, \bar{W}_1 \bar{W}_2 W_3, \ldots)$$

where W_i are independent beta variables with parameters (β, ρ), and $\bar{W}_i = 1 - W_i$.

2. *For $\beta = 0$, as $n \to \infty$,*

$$E[D_n(\rho, \beta)] \sim \log^2(n)$$

The above relation holds asymptotically for the gamma process as well. Moreover, the expected number of particles in the first block, $E[\#B_{1,n}]$, is $n/(\rho \log n)$. Asymptotically,

$$\frac{\log \#B_{1,n}}{\log n} \xrightarrow{D} U$$

where U has the uniform distribution on $(0, 1)$.

3. *For $\beta \in (-1, 0)$, as $n \to \infty$,*

$$E[D_n(\rho, \beta)] \sim n^{-\beta}$$

The fraction of edges in the first block is asymptotically n^{β}. Asymptotically, for all $\rho > 0$,

$$\mathrm{pr}(\#B_{1,n} = d) \to g_\beta(d)$$

where $g_\beta(d)$ is defined by

$$g_\beta(d) = \frac{-\beta \cdot \Gamma(d + \beta)}{\Gamma(d + 1) \cdot \Gamma(1 + \beta)} = \frac{-\beta}{\Gamma(1 + \beta)} d^{\beta - 1}$$

for large d. So the number of particles in the first block has a power law distribution of degree $1 - \beta$.

The proof can be found in Appendix C.

3 Exchangeable mixture model

3.1 Introduction

The development in this section follows closely Kalbfleisch (1978), Clayton (1991), and Hjort (1990). It considers survival processes driven by an arbirary, completely independent, stationary random measure, not necessarily the gamma process. This section considers only exchangeable processes, but these methods can be easily extended to handle inhomogeneity.

3.2 Completely independent random measure

A non-negative measure Λ on $\mathbb{R}$ is said to be *completely independent* if the random variables $\Lambda(A_1), \ldots, \Lambda(A_n)$ are independent whenever $A_1, \ldots, A_n$ are disjoint subsets (Kingman 1993, chapter 8), and *stationary* if the distribution is unaffected by translation in $\mathbb{R}$. The distribution of $\Lambda(A)$ is necessarily infinitely divisible, and the characteristic exponent

$$\log E(e^{-t\Lambda(A)}) = -\nu(A)\,\Psi(t) \tag{3.1}$$

for $t \geq 0$ is essentially the same as the cumulant generating function. The characteristic exponent need not be analytic at the origin, so, even if the cumulants are finite, there need not be a Taylor expansion to generate them. The measure is stationary if $\nu(A) = \nu\,|A|$ is proportional to Lebesgue measure on $\mathbb{R}$, which is subsequently assumed for simplicity of notation (i.e. $\nu \geq 0$, a constant). The Lévy-Khintchine characterization for positive random variables implies

$$\Psi(t) = \gamma t + \int_0^\infty (1 - e^{-zt})\,dw(z)$$

for some $\gamma \geq 0$ and some measure w on $(0, \infty]$, called the Lévy measure, such that the integral is finite for $t > 0$.

The random measure has a simple interpretation in terms of a Poisson point process $X \subset \mathbb{R} \times (0, \infty)$ with Lévy intensity $dt\,dw(z)$ such that

$$\Lambda(A) = \gamma|A| + \int_{A\times(0,\infty)} z\,dX$$

is Lebesgue measure plus the sum of the z-components of the points of X that fall in $A \times (0, \infty)$. If the Lebesgue component is missing, i.e. $\gamma = 0$, then Λ has countable support. If the Lévy measure is finite, then Λ has finitely many atoms in bounded subsets; otherwise, if w is not finite, Λ has countable dense support.

3.3 Survival process

Let Λ be a stationary, completely independent, random measure on $\mathbb{R}$, and let $T_1, T_2, \ldots$ be conditionally independent and identically distributed random variables such that

$$\mathrm{pr}(T_i > t \,|\, \Lambda) = \exp(-\Lambda(0, t])$$

for $t \geq 0$. In other words, the survival times are non-negative with cumulative conditional hazard $H(t) = \Lambda(0, t]$. Conditional on Λ, the multivariate survival function is

$$\mathrm{pr}(T_1 > t_1, \ldots, T_n > t_n \,|\, \Lambda) = \exp\Big(-\int_0^\infty R^\sharp(s)\,d\Lambda(s)\Big),$$

where $R^\sharp(s) = \#\{i\colon t_i \geq s\}$. From the definition (3.1) of the characteristic exponent, the unconditional joint survival function is

$$\operatorname{pr}(T_1 > t_1, \ldots, T_n > t_n) = \exp\Big(-\nu \int_0^\infty \Psi(R^\sharp(s))\, ds\Big), \tag{3.2}$$

depending on the the characteristic exponent at integer values. The following connects the survival process associated with completely independent random measures to the set of Markov survival processes.

Theorem 3.1. *The unconditional risk set evolves as a Markov process with characteristic index $\zeta_n = \nu\Psi(n)$ evaluated at the positive integers.*

Proof. Given Λ, the conditional density at the risk-set trajectory R consisting of $k \leq n$ failure times is

$$\exp\Big(-\int_0^\infty R^\sharp(t)\, d\Lambda\Big) \prod_{j=1}^k \big(1 - e^{-\Lambda(dt_j)}\big)^{d_j}$$

where $d_j \geq 1$ is the number of failures at t_j. Bearing in mind that Λ is completely independent and that $E(e^{-s\Lambda(A)}) = e^{-\nu|A|\Psi(s)}$, the contribution to the unconditional density of one failure time with $R^\sharp(t^-) = d + r$ and $R^\sharp(t) = r$ is an alternating sum

$$\begin{aligned} E\Big(e^{-r\Lambda(dt)}\, \big(1 - e^{-\Lambda(dt)}\big)^d\Big) &= \sum_{j=0}^d (-1)^j \binom{d}{j} e^{-\nu\, dt\, \Psi(r+j)} \\ &= \nu\, dt \sum_{j=0}^d (-1)^{j+1} \binom{d}{j} \Psi(r+j) + o(dt) \\ &= \nu\, (-1)^{d-1} (\Delta^d \Psi)(r)\, dt + o(dt). \end{aligned}$$

Note that $(-1)^{d-1}(\Delta^d\Psi)(r) \geq 0$ and is decreasing in both arguments. It follows that the joint density at R is

$$f_n(R) = \nu^k \exp\Big(-\nu \int_0^\infty \Psi\big(R^\sharp(s)\big)\, ds\Big) \times \prod_{j=1}^k (-1)^{d_j - 1} (\Delta^{d_j}\Psi)(r_j)\, dt_j \tag{3.3}$$

which is the same as (2.9). □

A followup question of interest is whether the converse is also true. Theorem 3.2 answers in the affirmative.

Theorem 3.2. *Every Markov survival process is generated by a completely independent random measure*

Proof. Recall that the set of singleton splitting rules, $\{q(n,1)\}$, completely determine the splitting rule by consistency.

Equation (2.3) provides an integral representation for the singleton splitting rule for all Markov survival processes.

$$q(n,1) \propto c + \int s^n(1-s)\varpi(ds)$$

Suppose there were an associated completely independent random measure. Then the Lévy-Khintchine characterization implies that the singleton splitting rule is

$$q(n,1) = \nu \cdot (\Psi(n+1) - \Psi(n)) \tag{3.4}$$

$$= \nu \cdot \left(\gamma + \int_0^\infty e^{-nz}(1-e^{-z})w(dz)\right) \tag{3.5}$$

Set $\gamma = c$, and $w(dz) = e^{-z}\varpi(-\log(s) \in dz)$ and without loss of generality assume $\nu = 1$, as this corresponds to the standardized form of the characteristic index. It rests to check that the measure w is a Lévy measure. It is easy to show that

$$\int_0^1 (1-s)\,\varpi(ds) < \infty \Rightarrow \int_0^\infty (1-e^{-zt})w(dz) < \infty, \quad t \geq 0$$

Therefore the Markov survival process is generated by a completely independent random measure. □

3.4 Example: homogeneous gamma process

The density of the gamma distribution $Z \sim \mathrm{Ga}(\nu,\rho)$ at $z > 0$ is

$$z^{\nu-1}\rho^\nu e^{-z\rho}/\Gamma(\nu).$$

For $\nu > 0, \rho > 0$, the distribution has finite moments of all orders, and moment generating function

$$M(t) = E(e^{tZ}) = \left(\frac{\rho}{\rho - t}\right)^\nu$$

$$K(t) = \log M(t) = -\nu \log\bigl(1 - t/\rho\bigr) = -\nu\Psi(-t)$$

for $t < \rho$. It follows that the distribution is infinitely divisible with rth cumulant $\nu\,\Gamma(r)/\rho^r$. It is convenient to extend the parameter space to include $\nu = \infty$, in which case $\mathrm{Ga}(\infty,\rho)(\{\infty\}) = 1$ for all $\rho > 0$, i.e., unit mass is assigned to the point at infinity.

The *homogeneous gamma process* with parameter (ν,ρ) on the real line is a stationary random measure Λ, which has the following properties:

- for each Borel subset $A \subset \mathbb{R}$ of Lebesgue measure $|A|$, the random variable $\Lambda(A)$ is distributed according to the gamma distribution $\mathrm{Ga}(\nu|A|,\rho)$;

- the values $\Lambda(A_1), \ldots, \Lambda(A_n)$ assigned by Λ to disjoint subsets $A_1, \ldots, A_n$ are independent random variables;
- with probability one, Λ is purely atomic with countable dense support.

To explain the reasoning behind the last point, let $\mathbf{X}$ be a Poisson process with intensity $\nu\, z^{-1} e^{-\rho z}\, dx\, dz$ at (x, z) in the product space $\mathbb{R} \times (0, \infty)$. Then

$$\Lambda(A) = \int_{A \times (0,\infty)} z\, d\mathbf{X}$$

is distributed as the gamma process with parameter (ν, ρ). In other words, to each $(x, z) \in \mathbf{X}$ there corresponds an atom of Λ such that $\Lambda(\{x\}) = z$. If $0 < |A| < \infty$, the set $\mathbf{X} \cap (A \times (0, \infty))$ is infinite, but the subset $\mathbf{X} \cap (A \times (\epsilon, \infty))$ is finite for every $\epsilon > 0$. The value assigned by Λ to $A \subset \mathbb{R}$ is the sum of the z-components of all points of $\mathbf{X}$ that lie in $A \times (0, \infty)$, and this sum is finite with probability one. Since there are no coincident atoms, i.e., no pairs of points in $\mathbf{X}$ whose x-components are equal, the Poisson process is equivalent to Λ. For further details, see Kingman (1993, chapter 8).

The characteristic index for the gamma process is $\Psi(t) = \log(1 + t/\rho)$, which means that

$$\begin{aligned}
\Delta \Psi(t) &= \log(\rho + t + 1) - \log(\rho + t) \\
(-1)\Delta^2 \Psi(t) &= -\log(\rho + t + 2) + 2\log(\rho + t + 1) - \log(\rho + t) \\
(-1)^{d-1} \Delta^d \Psi(t) &= \sum_{j=0}^{d} (-1)^{j-1} \binom{d}{j} \log(\rho + t + j) \\
&\simeq \Gamma(d) \;/\; (\rho + t + 1/2)^{\uparrow d}. \qquad (3.6)
\end{aligned}$$

The joint marginal density of the survival times in the gamma process is obtained by substitution into (3.3):

$$\nu^k \exp\!\left(-\nu \int_0^\infty \log\bigl(1 + R^\sharp(s)/\rho\bigr)\, ds \right) \times \prod_{j=1}^{k} (-1)^{d_j - 1} (\Delta^{d_j} \log)(\rho + r_j)\, dt_j$$

Approximation (3.6), which holds with relative error $O(1/t^2)$ for large t, is the dth order forward difference of the characteristic exponent $\psi(\rho' + t) - \psi(\rho')$, of the harmonic process with parameter $\rho' = \rho + 1/2$. The discrete component is thus approximately a product over failure times

$$\prod_{j=1}^{k} \frac{\Gamma(d_j)}{(\rho' + r_j)^{\uparrow d_j}}$$

where d_j is the number of failures and $r_j = R^\sharp(t_j)$ is the size of the risk set at time t_j. In particular, if there are no censored records, the joint density for the

harmonic process with parameter ν, ρ is

$$\frac{\nu^k}{\rho^{\uparrow n}} \exp\biggl(-\nu \int_0^\infty \bigl(\psi(\rho + R^\sharp(s)) - \psi(\rho)\bigr)\, ds\biggr) \prod_{j=1}^k \Gamma(d_j),$$

where $\sum d_j = n$ and the integral in the exponent may be written as a finite sum

$$\sum_{k=0}^{n-1} \frac{t_{(n-k)}}{k+\rho}$$

over the ordered survival times $t_{(1)} \le t_{(2)} \le \cdots \le t_{(n)}$.

3.5 Parameter estimation

Consider the problem of parameter estimation for a two-parameter Markov survival process with characteristic index of the form $\zeta_n = \nu\Psi(n)$ where $\Psi(n) = \Phi(n+\rho) - \Phi(\rho)$ for $\nu > 0$, $\rho > 0$ with $\Phi(\cdot)$ given. Such a family is generated from a family of Lévy measures proportional to $w(dz)e^{-\rho z}$, so Λ is expected to have larger atoms if ρ is small. The gamma and harmonic processes are of this form with $w((0,1))$ and $w((1,\infty))$ both infinite.

The first goal is to estimate the parameters ρ, ν, from observations $T_1, \ldots, T_n$, some of which may be right censored. Consistency as $n \to \infty$ is not to be expected, nor is it necessarily important. Parameter estimation is usually an intermediate step, which is needed primarily to compute the Bayes estimate, or empirical Bayes estimate, of the survival distribution $F(t) = \mathrm{pr}(T_{n+1} > t \,|\, \mathrm{data})$.

For fixed ρ, the survival model (3.3) is a one-parameter family with a two-dimensional sufficient statistic

$$k, \quad \int_0^\infty \Psi(R^\sharp(t))\, dt,$$

where $1 \le k \le n$ is the number of blocks, or more generally the number of distinct death times. The first derivative of the log likelihood with respect to $\log \nu$ is $k - \nu \int \Psi(R^\sharp(t))\, dt$, the maximum-likelihood estimate is the ratio

$$\hat\nu^{-1} = \frac{1}{k} \int_0^\infty \Psi(R^\sharp(t))\, dt. \tag{3.7}$$

The Fisher information for $\log \nu$ is $E(k)$, suggesting that the asymptotic variance of $\log \hat\nu$ is $1/E(k)$. For the gamma and harmonic processes, $\log(n) < E(k) < n^\epsilon$ implies that the estimator is consistent in the absence of censoring. but the rate of convergence is very slow.

Estimation of ρ by maximum likelihood is certainly feasible, but a little more difficult. One natural option is to consider the product of the per-particle death rate $\nu\Psi(1)$ and the total particle time at risk $\int R^\sharp(t)\, dt$, and to estimate ρ by setting the product to the observed number of deaths, i.e.,

$$\#\mathrm{deaths} = \hat\nu \hat\Psi(1) \int_0^\infty R^\sharp(t)\, dt. \tag{3.8}$$

In the absence of censoring, this is equivalent to setting the mean survival time $\bar{T}_n$ to its expected value $1/\hat{\zeta}_1$. But $\mathrm{cov}(T_i, T_j) = 1/\zeta_2^2$ for each pair $i \neq j$ implies that

$$\mathrm{var}(\bar{T}_n) = 1/(n\zeta_1^2) + (n-1)/(n\zeta_2^2) \to 1/\zeta_2^2$$

does not tend to zero as $n \to \infty$. Nonetheless, this second equation is less sensitive than the first to rounding of survival times, which is a desirable property for applied work. The parameter pair can be estimated by iteration.

3.6 Numerical example

Consider parameter estimation for a set of failure and censoring times (in weeks) of the 6-MP subset of leukemia patients taken from Gehan (1965):

$$6, 6, 6, 6^\star, 7, 9^\star, 10, 10^\star, 11^\star, 13, 16, 17^\star, 19^\star, 20^\star, 22, 23, 25^\star, 32^\star, 32^\star, 34^\star, 35^\star$$

There are 9 uncensored observations, and a total risk time of 359 weeks. Assuming the survival times are iid exponential with rate parameter, θ, then the maximum likelihood estimate of θ is given by 9/359, or an expected survival time of 39.89 weeks.

Consider the two-parameter Markov survival process defined in section 3.5, specifically the harmonic and gamma processes. Table 1 provides maximum likelihood estimates for ρ and ν. For the gamma process, the empirical Bayes estimate of the rate is then $\hat{\nu} \cdot \log(1 + \hat{\rho}^{-1}) \approx 2.47 \times 10^{-2}$, implying the expected survival time is 40.52 weeks. The expected time is the same for the harmonic process.

Table 1: Maximum likelihood estimates for two processes

	Harmonic process		Gamma process	
Parameter	Est.	Std. Error	Est.	Std. Error
ρ	21.45	19.63	20.95	19.61
ν	0.53	0.44	0.53	0.44

Estimation using the maximum likelihood estimate of ν given ρ and the natural relation between the marginal survival rate associated with the gamma process and $\hat{\theta}$

$$\hat{\nu}(\rho)\,(\Phi(1+\rho) - \Phi(\rho)) = \hat{\theta} = 9/359$$

yields $(\hat{\rho}, \hat{\nu}) = (19.24, 0.49)$ for the gamma process and $(19.73, 0.49)$ for the harmonic process. Supplementary figure 3 shows that the profile likelihood for ρ is relatively flat for values sufficiently removed from the origin. For the harmonic process, the figure suggests a 95% confidence interval of approximately $[1.3, 5.1]$ for $\log(\rho)$, while under the gamma process, there is an approximate confidence interval of $[1.2, 5.1]$. Twice the difference in the log-likelihoods at their respective maxima is 8.12×10^{-5}.

Figure 1 shows the conditional survival distribution given the observed risk set trajectory. The empirical Bayes estimate of the conditional distribution

for the harmonic process is approximately equal to that of the gamma process. Both are approximately an average of the Kaplan-Meier product limit estimator and the maximum likelihood exponential estimator of the conditional survival distribution.

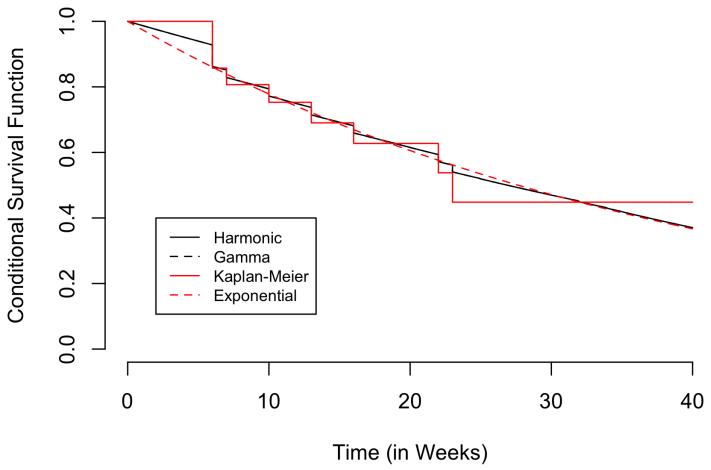

Figure 1: Conditional survival distribution for leukemia patients

4 Ties as a result of numerical rounding

Up to this point, individuals having the same recorded survival time are regarded as failing simultaneously. Now consider the case where the actual failure times are distinct, so that tied values arise solely as a result of numerical rounding. The integral in the exponent is a continuous function of the risk set trajectory $R(t)$, so an ϵ-perturbation of failure times has an $O(\epsilon)$ effect on the integral, which is ignored. However, the remaining term is not a continuous function of the observations, so an ϵ-perturbation by rounding may have an appreciable effect on the likelihood. Most obviously, the statistic k, the number of distinct failure times, is not continuous as a function of T; if ties are an artifact of rounding, then k is the total number of failures.

While the likelihood and parameter estimation are affected by ties as a result of numerical rounding, the conditional survival distribution for the harmonic process given ρ and ν is unaffected due to the weak continuity of predictive distributions. This suggests it may be best to regard ρ as a fixed "tuning parameter". As all other processes have discontinuous predictive distributions, the use of the harmonic process in applications where ties are likely the result of numerical rounding seems most natural.

5 Conclusion

Presented here is the set of exchangeable, consistent Markov survival processes defined by their characteristic index, ζ. These processes are in correspondence with the set of completely independent random measures, each determined by the characteristic exponent.

Markov survival processes exhibit multiple tied failure times, generating exchangeable partial rankings among individuals. It has been shown that the number of blocks depends on the characteristic index. The harmonic and gamma processes are studied in detail, showing the number of blocks grows asymptotically at a rate of $\log^2(n)$. Censoring is easily incorporated and parameter estimation given a set of survival and censoring times is examined. The impact of ties as the result of numerical rounding concludes the discussion.

References

[1] Aldous, D. (1996) Probability distributions on cladograms. In *Random Discrete Structures* **76**, 1–18.

[2] Breslow, N. (1974) Covariance analysis of censored survival data. *Biometrics* **30**, 89–99.

[3] Clayton, D.G. (1991) A Monte Carlo method for Bayesian inference in frailty models. *Biometrics* **47**, 467–485.

[4] Cornfield, J., and Detre, K. (1977) Bayesian life table analysis. *J. Roy. Statist. Soc. B* **39**, 264–296.

[5] Cox, D.R. and Oakes, D. (1978) *Analysis of Survival Data.* London, Chapman & Hall.

[6] Doksum, K.A. (1974) Tailfree and neutral random probabilities and their posterior distributions. *Annals of Probability* **2**, 183–201.

[7] Efron, B. (1975) The efficiency of Cox's likelihood function for censored data. Technical Report.

[8] Ferguson, T.S. (1973) A Bayesian analysis of some nonparametric problems. *Annals of Statistics* **1**, 209–230.

[9] Gehan, E.A. (1965) A generalized Wilcoxon test for comparing arbitrarily single-censored samples. *Biometrika* **52**, 203–233.

[10] Haas, B. , Miermont, G., Pitman, J., and Winkel, M. (2008) Continuum tree asymptotics of discrete fragmentations and applications to phylogentic models. *Annals of Probability* **36**, 1790–1837.

[11] Hjort, N.L. (1990) Nonparametric Bayes estimators based on beta processes in models for life history data. *Annals of Statistics* **18**, 1259–1294.

[12] James, L.F. (2007) Neutral-to-the-right species sampling mixture models. Chapter 21 in *Advances in Statistical Modeling and Inference.* Essays in honor of K. Doksum. 425–439, World Scientific Series in Biostatistics, **3**.

[13] Kalbfleisch, J.D. (1978) Nonparametric Bayesian analysis of survival time data. *J. Roy. Statist. Soc. B* **40**, 214–221.

[14] Kaplan, E. L. and Meier, P. (1958) Nonparametric estimation from incomplete observations. *J. Amer. Statist. Assn.* **53**, 457-481.

[15] Kingman, J.F.C (1993) *Poisson Processes.* Oxford Scientific Publications.

[16] Kingman, J.F.C (1978). The representation of partition structures. *J. London Math. Soc.* **2**, 374–380.

[17] McCullagh, P., Pitman, J., and Winkel, M. (2008) Gibbs fragmentation trees. *Bernoulli* **14**, 988–1002.

[18] Müller, P., and Quintana, F. (2004) Nonparametric Bayesian data analysis. *Statistical Science* **19**, 95–110.

[19] Pitman, J. (2006) *Combinatorial Stochastic Processes.* Lecture Notes Math. Springer, New York.

[20] Walker, S., Damien, P., Laud, P., and Smith, A. (1999) Bayesian nonparametric inference for random functions and related functions. *J. Roy. Statist. Soc. B* **61**, 485–527.

A Supplementary Figures

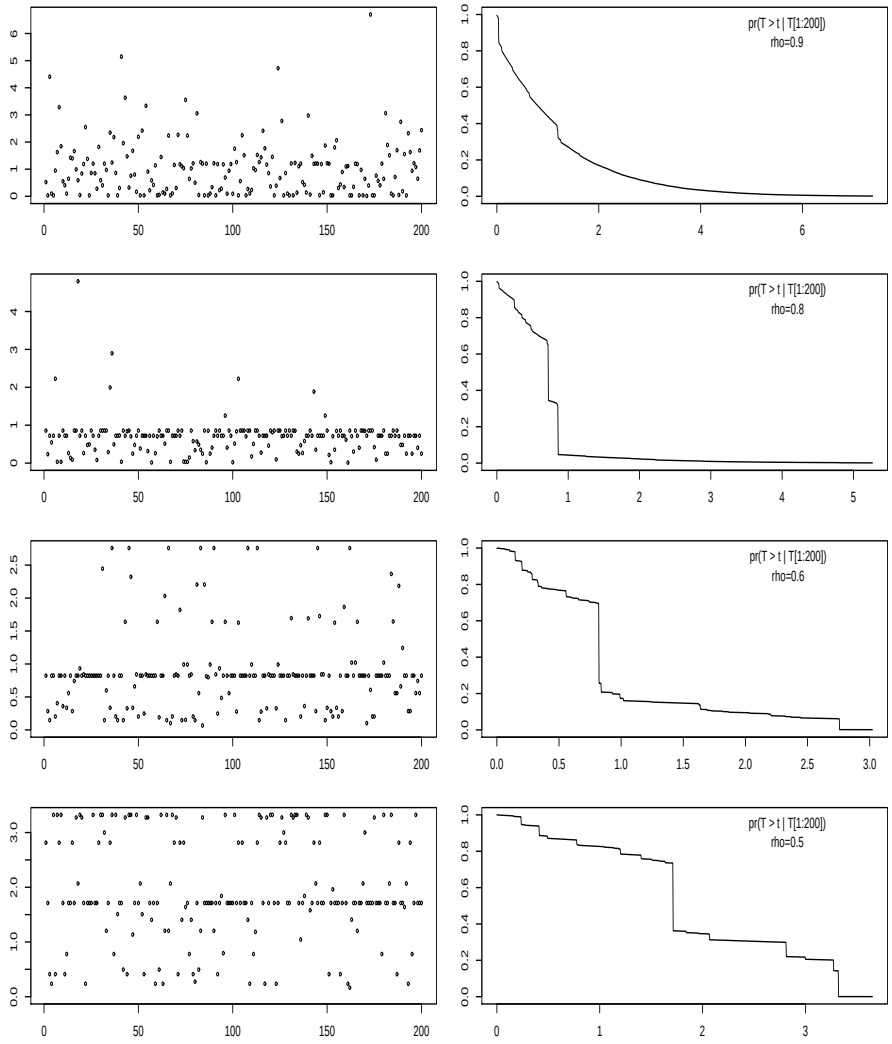

Supplementary Figure 1: Simulated survival process with index $\zeta_n \propto n^\rho$, together with the conditional survival function (right panel). Parameter values: $\rho = 0.9, 0.8, 0.6, 0.5$. Among the first 200 failure times, the number of distinct values was 150, 65, 42, 25, respectively, but these are highly variable from run to run.

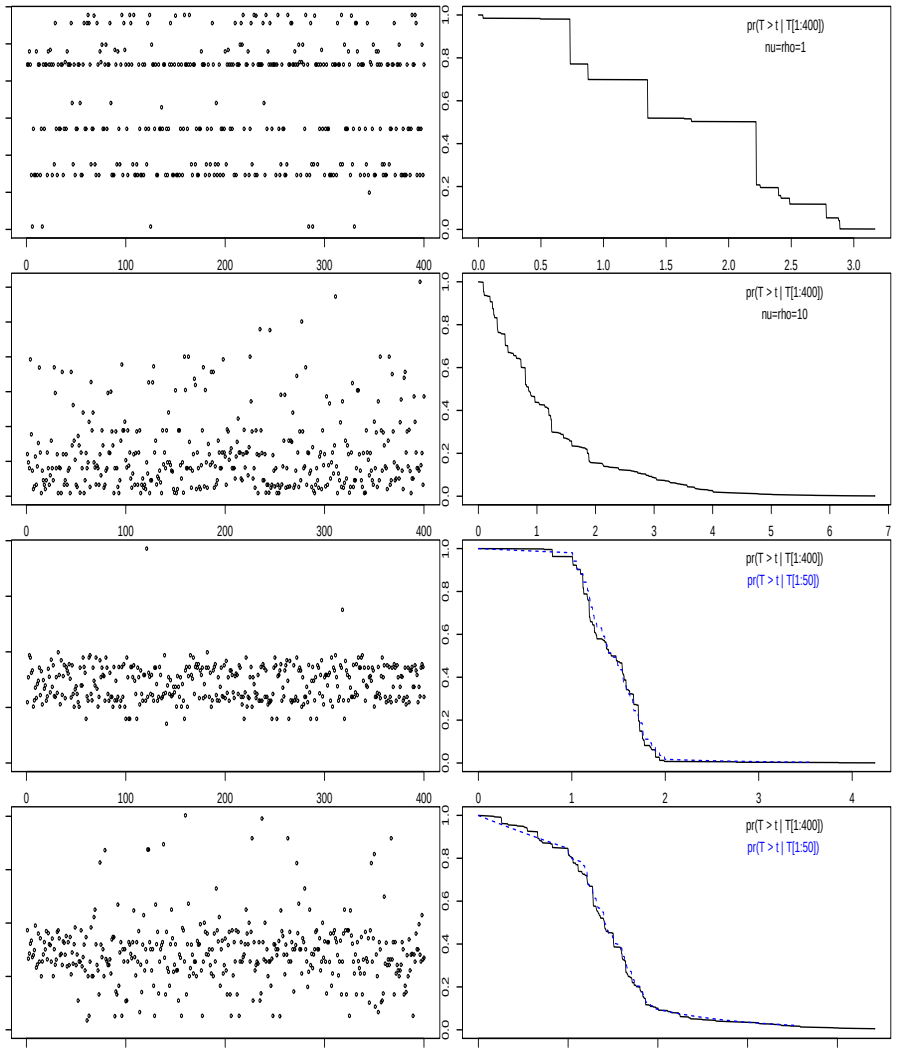

Supplementary Figure 2: Simulated harmonic survival process, together with the conditional survival function (right panel). The lower two series are seeded with 50 initial values, independent, uniform on $(1, 2)$. Parameter values: $\nu = \rho = 1$ in rows 1, 3; $\nu = \rho = 10$ in rows 2, 4. Among the first 400 failure times, the number of distinct values was 15, 93, 58, 110, respectively, including the initial seeds.

Supplementary Figure 3: Profile Log-Likelihood of $\log(\rho)$

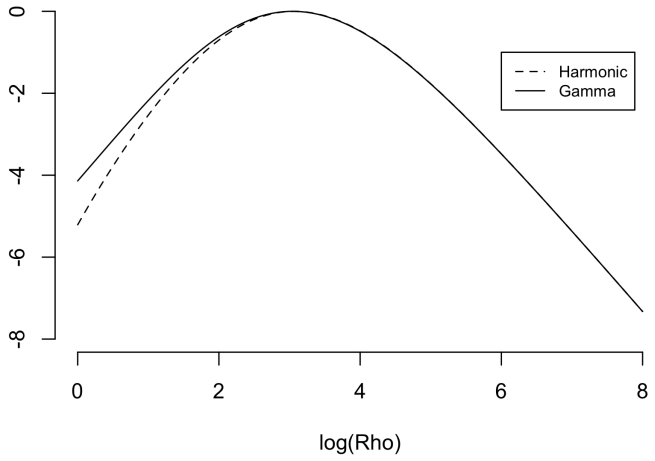

Supplementary Figure 4: Profile log-likelihood of β_1

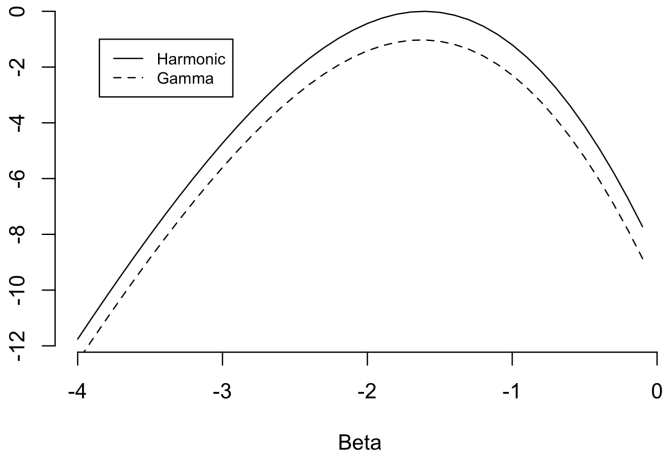

B Proof of Proposition 2.6

Proof. A consistent splitting rule is completely determined by the set of singleton splitting rules, $\{q(n,1) \,|\, n \geq 0\}$ where $q(0,1) = 1$ by definition. It is immediately apparent from the construction that $q_{r,1} = \Delta\zeta_r/\zeta_{r+1} = 1 - \zeta_r/\zeta_{r+1}$. Considering the standardized sequence and using this relation yields a one-to-one correspondence between the set of singleton splitting rules and the characteristic index. Therefore the characteristic index completely determines a splitting rule.

Showing equation (2.5) yields a splitting rule that satisfies the consistency condition (2.2) completes the proof. The left hand side is given by

$$\begin{aligned}(1 - q(n,1))q(n-d,d) &= \frac{\zeta_n}{\zeta_{n+1}} \frac{(-1)^{d-1}(\Delta^d\zeta)_{n-d}}{\zeta_n} \\ &= \frac{(-1)^{d-1}(\Delta^d\zeta)_{n-d}}{\zeta_{n+1}}\end{aligned}$$

The right hand side is given by

$$\begin{aligned}q(n-d+1,d) + q(n-d,d+1) &= \frac{(-1)^{d-1}(\Delta^d\zeta)_{n-d+1}}{\zeta_{n+1}} + \frac{(-1)^{d}(\Delta^{d+1}\zeta)_{n-d}}{\zeta_{n+1}} \\ &= \frac{(-1)^{d-1}}{\zeta_{n+1}} \cdot \left[(\Delta^d\zeta)_{n-d+1} - (\Delta^{d+1}\zeta)_{n-d}\right] \\ &= \frac{(-1)^{d-1}}{\zeta_{n+1}} \cdot (\Delta^d\zeta)_{n-d}\end{aligned}$$

So the splitting rule automatically satisfies the consistency condition. By definition, the splitting rule must be non-negative, and therefore the sequence defines a consistent splitting rule if the forward differences are non-negative, $(-1)^{d-1}(\Delta^d\zeta)_r \geq 0$ for all $r \geq 0$, $d \geq 1$. □

C Proof of Theorem 2.9

Proposition C.1. *Define $D_n(\rho)$ denote the number of blocks given n individuals and $\rho \in (0,\infty)$. As $n \to \infty$,*

$$E[D_n(\rho)] \sim \log^2(n)$$

for both the harmonic and gamma processes, where the ratio tends to a constant equal to $\frac{1}{2\cdot\Psi_1(\rho)}$ where Ψ_1 is the trigamma function.

Proof. For the gamma process, the recurrence relation of the expected number

of blocks is given by:

$$\begin{aligned}\mu_n = E[D_n] &= 1 + \sum_{d=1}^{n} \binom{n}{d} q(n-d,d) \cdot \mu_{n-d} \\ &= 1 + \sum_{d=1}^{n} \binom{n}{d} \frac{(-1)^{d-1}(\Delta^d \zeta)_{n-d}}{\log(1+\frac{n}{\rho})} \cdot \mu_{n-d}\end{aligned}$$

which gives the following approximation:

$$\begin{aligned}(-1)^{d-1}(\Delta^d \zeta)_{n-d} &= \sum_{i=0}^{d-1} \binom{d-1}{i} (-1)^i \log\left(1 + \frac{1}{R+i+\rho}\right) \\ &\approx \sum_{i=0}^{d-1} \binom{d-1}{i} (-1)^i \frac{1}{R+i+\rho} \\ &= \frac{\Gamma(d)\Gamma(n-d+\rho)}{\Gamma(n+\rho)} = \frac{\Gamma(d)}{(n-d+\rho)^{\uparrow d}}\end{aligned}$$

This is the exact expression for the harmonic process. So for both processes, the above implies that the probability of d individuals in block 1 is proportional to:

$$\begin{aligned}\binom{n}{d}(\Delta^d \zeta)_{n-d} &= \frac{\Gamma(n+1)}{\Gamma(n+\rho)} \cdot \frac{\Gamma(n-d+\rho)}{\Gamma(n-d+1)} \cdot \frac{1}{d} \\ &\approx \left(\frac{n-d}{n}\right)^{\rho-1} \cdot \frac{1}{d}\end{aligned}$$

where the approximation

$$\frac{\Gamma(x+1)}{\Gamma(x+\rho)} \approx x^{\rho-1}$$

is used. Substition of this into the recurrence relation yields:

$$0 \approx 1 + \frac{1}{\log\left(1+\frac{n}{\rho}\right)} \sum_{d=1}^{n} \left(\frac{n-d}{n}\right)^{\rho-1} \cdot \frac{1}{d} \left(\mu_{n-d} - \mu_n\right)$$

Writing $\mu_n \sim c \log^2(n)$ and approximating the sum by an integral yields:

$$\begin{aligned}0 &\approx 1 + \frac{c}{\log\left(1+\frac{n}{\rho}\right)} \int_0^1 \frac{(1-x)^{\rho-1}\log(1-x)}{x} \left(2\log(n) - \log(1-x)\right) dx \\ &\to 1 + 2 \cdot c \cdot \int_0^1 \frac{(1-x)^{\rho-1}\log(1-x)}{x} dx \\ &= 1 - 2 \cdot c \cdot \Psi_1(\rho)\end{aligned}$$

Giving us $c = \frac{1}{2 \cdot \Psi_1(\rho)}$ as required. □

A similar argument can be used to obtain the asymptotic expected fraction of particles in the first block.

Proposition C.2. *The expected number of particles in the first block, $E[\#B_1]$, is $n/(\rho \log n)$ for the harmonic and gamma process. Asymptotically, for integer values of ρ,*

$$\frac{\log B_1}{\log n} \xrightarrow{D} U$$

where U has the uniform distribution on $(0, 1)$.

Proof. The expected number of particles in the first block is given by

$$\begin{aligned}
\sum_{d=1}^{n} d\binom{n}{d} q(n-d, d) &= \frac{1}{\log(1+\frac{n}{\rho})} \sum_{d=1}^{n} d \left(\frac{n-d}{n}\right)^{\rho-1} \frac{1}{d} \\
&\to \frac{n}{\log(1+\frac{n}{\rho})} \int (1-x)^{\rho-1}\, dx \\
&\approx \frac{n}{\log(n)} \frac{1}{\rho}
\end{aligned}$$

So the fraction of particles in the first block is roughly $(\rho \log(n))^{-1}$. For $z \in (0, 1)$, the asymptotic distribution is given by

$$\begin{aligned}
\text{pr}\left[\frac{\log B_1}{\log n} \leq z\right] &= \frac{1}{\log(n)} \int_{1/n}^{n^{z-1}} x^{-1}(1-x)^{\rho-1} dx \\
&= \frac{1}{\log n} \left[\sum_{j=0}^{p-1} c_j x^j + \log(x)\right]_{1/n}^{n^{z-1}} \\
&\to \frac{1}{\log n} \left[(z-1)\log(n) - \log(1/n))\right] = z
\end{aligned}$$

where the second line is true for $\rho \in \mathbb{Z}$. The result holds for general $\rho > 0$ by the squeezing theorem. □

We now consider the number of blocks and block sizes for the general beta-splitting rules.

Proposition C.3. *Define $D_n(\rho, \beta)$ denote the number of blocks given n individuals for the beta process. As $n \to \infty$,*

$$E[D_n(\rho, \beta)] \sim \log(n)$$

where the ratio tends to a constant $c = \frac{1}{\psi(\rho+\beta)-\psi(\rho))}$. The fraction of edges in the first block, $\#B_1$, is distributed $Beta(\rho, \beta)$. Therefore, the relative frequencies within each block is given by

$$(P_1, P_2, \ldots) = (B_1, \bar{B}_1 B_2, \bar{B}_1 \bar{B}_2 B_3, \ldots)$$

where B_i are independent beta variables with parameters (ρ, β).

Proof. The probability of d individuals in the first block is given by

$$\binom{n}{d} q(n-d,d) \approx \frac{1}{Z_n}(1-x)^{\rho-1}x^{\beta-1}$$

where $Z_n = \int_0^1 (1-s^n)s^{\rho-1}(1-s)^{\beta-1}ds = B(\rho,\beta) - B(\rho+n,\beta)$. As $n \to \infty$, for $\rho, \beta > 0$ the normalization constant converges to $B(\rho,\beta)$. Therefore, the expected number of blocks is given by

$$0 = 1 + \frac{1}{Z_n}\sum_{d=1}^{n}\left(\frac{n-d}{n}\right)^{\rho-1}\left(\frac{d}{n}\right)^{\beta-1} n^{-1}\left(\mu_{n-d} - \mu_n\right)$$

Writing $\mu_n \sim c\log(n)$ then approximating the sum by an integral yields

$$\begin{aligned} 0 &\approx 1 + \frac{c}{B(\rho,\beta)}\int_0^1 (1-x)^{\rho-1}x^{\beta-1}\log(1-x)dx \\ &= 1 + c\cdot E[\log(1-X)] \\ &= 1 + c\cdot(\psi(\rho) - \psi(\rho+\beta)) \end{aligned}$$

as desired. The probability of d out of n particles in block one is given by

$$\frac{n}{Z_n}\left(\frac{n-d}{n}\right)^{\rho-1}\left(\frac{d}{n}\right)^{\beta-1} n^{-1} \to \frac{n}{B(\rho,\beta)}(1-x)^{\rho-1}x^{\beta-1}dx$$

So the fraction of particles in block one is distributed beta with parameters (ρ,β). □

The final case is when $\beta \in (-1,0)$. Then the number of blocks grows polynomially in n.

Proposition C.4. *Define $D_n(\rho,\beta)$ denote the number of blocks given n individuals for the beta process. For $\beta \in (-1,0)$, as $n \to \infty$,*

$$E[D_n(\rho,\beta)] \sim n^{-\beta}$$

where the ratio tends to a constant $c = -\Gamma(\rho+\beta+1)/(\Gamma(\rho)\beta^2)$. The fraction of edges in the first block is asymptotically proportional to n^β. Asympotically,

$$\mathrm{pr}[B_1 = d] = \frac{-\beta\cdot\Gamma(d+\beta)}{\Gamma(d+1)\cdot\Gamma(1+\beta)} = \frac{-\beta}{\Gamma(1+\beta)}d^{\beta-1}$$

for large d.

Proof. For $\beta \in (-1, 0)$, the characteristic index is given by

$$\begin{aligned}
\zeta_n &= \int_0^1 (1 - s^n) s^{\rho-1} (1-s)^{\beta-1} ds \\
&= \int_0^1 \sum_{j=0}^{n-1} s^{j+\rho-1} (1-s)^{\beta} ds \\
&= \sum_{j=0}^{n-1} \frac{\Gamma(j+\rho)\Gamma(\beta+1)}{\Gamma(j+1+\beta+\rho)} \\
&\approx \frac{\Gamma(1+\beta)}{-\beta} n^{-\beta}
\end{aligned}$$

for n large. Plugging this into the recursive formula, while assuming $\mu_n \sim c n^{-\beta}$ yields

$$\begin{aligned}
0 &= 1 + c \frac{-\beta\, n^{\beta}}{\Gamma(1+\beta)} \sum_{d=1}^{n} \left(\frac{n-d}{n}\right)^{\rho-1} \left(\frac{d}{n}\right)^{\beta-1} n^{-1} \left((n-d)^{-\beta} - n^{-\beta}\right) \\
&= 1 - \frac{c\,\beta}{\Gamma(1+\beta)} \sum_{d=1}^{n} \left(\frac{n-d}{n}\right)^{\rho-1} \left(\frac{d}{n}\right)^{\beta-1} n^{-1} \left((1-\frac{d}{n})^{-\beta} - 1\right) \\
&= 1 - \frac{c\,\beta^2}{\Gamma(1+\beta)} \sum_{d=1}^{n} \left(\frac{n-d}{n}\right)^{\rho-1} \left(\frac{d}{n}\right)^{\beta} n^{-1} \\
&\to 1 - \frac{c\,\beta^2}{\Gamma(1+\beta)} \int (1-x)^{\rho-1} (x)^{\beta}\, dx \\
&= 1 + c \frac{\beta^2\, \Gamma(\rho)}{\Gamma(\rho+\beta+1)}
\end{aligned}$$

where for large n the approximation $(1 - \frac{d}{n})^{-\beta} = 1 + \beta \frac{d}{n}$ is used. The result is immediate. The fraction of edges in the first block is given by

$$\frac{-\beta\, n^{\beta}}{\Gamma(1+\beta)} \sum_{i=1}^{n} \left(\frac{d}{n}\right)^{\beta} \left(\frac{n-d}{n}\right)^{\rho-1} n^{-1} = \frac{-\beta\, n^{\beta}}{\Gamma(1+\beta)} \int (1-x)^{\rho-1} (x)^{\beta}\, dx = \frac{-\beta\, \Gamma(\rho)}{\Gamma(\rho+\beta+1)} n^{\beta}$$

Finally, the asymptotic distribution comes from

$$\begin{aligned}
\mathrm{pr}(B_1 = d) &= \zeta_n^{-1} \frac{\Gamma(d+\beta)\Gamma(n-d+\rho)}{\Gamma(n+\rho+\beta)} \cdot \binom{n}{d} \\
&\sim \frac{-\beta \cdot n^{\beta}}{\Gamma(1+\beta)} \cdot \frac{\Gamma(n+1)}{\Gamma(n+\rho+\beta)} \frac{\Gamma(d+\beta)}{\Gamma(d+1)} \frac{\Gamma(n-d+\rho)}{\Gamma(n-d+1)} \\
&\sim \frac{-\beta}{\Gamma(1+\beta)} \frac{\Gamma(d+\beta)}{\Gamma(d+1)} \left(1 - \frac{d}{n}\right)^{\rho-1} \\
&\to \frac{-\beta}{\Gamma(1+\beta)} \frac{\Gamma(d+\beta)}{\Gamma(d+1)}
\end{aligned}$$

as $n \to \infty$. □